

NCC98 Network Backbone Design Project

**Backbone Infrastructure
and
High-End Switch
Evaluation and
Recommendation
Report**

15 December 1997

NCC98 Network Backbone Design Project

This document was prepared as an evaluation and recommendation report to the Network Control Center (NCC) Project Office, Code 530.5, at Goddard Space Flight Center for use in the NCC98 Project. It is part of a set of informal documents which were determined to be required by Task Management (GSA Manager) for this project (NCC98 Network Backbone Design Project). This document and any others distributed to date were prepared by:

Jeff Crowder
Systems Integration Engineer
Lockheed Martin Space Mission Systems & Services
7700 Hubble Drive
Seabrook, MD 20706
email: jcrowder@pop500.gsfc.nasa.gov
voice: 301-286-8975

Any questions or correspondence relative to information contained within this document should be directed to the author above.

Introduction

Purpose:

This report provides a recommendation for the basic hardware components of the backbone design for NCC98. It is the culmination of a year and a half long effort involving extensive research and testing in both the Building 12 N2 Lab and within the Development LAN (DEVLAN). The focus of the report details the evaluation of several high-end LAN switches and their respective manufacturers for potential use in the NCC98 LAN backbone. The report concludes with a recommendation for the switch which best meets the criteria established before and during the evaluation and which best fits the design of the NCC98 Network Backbone (refer to the *Hardware Specifications for the NCC98 Network Backbone Architecture* document as a baseline) and its respective applications. Additional recommendations are made concerning associated LAN infrastructure such as cabling, network interface cards (NICs) and shared LAN hubs.

Scope:

The scope of this report is limited to the NCC98 Network Backbone only which includes the LAN cabling, hubs, switches, network interface cards (NICs) and any other LAN infrastructure related components. The NCC98 Network Backbone includes the Test & Training (T&T) environment, the Auxiliary NCC (ANCC) environment, the Operational (OPS) environment, and the external network circuits to Nascom's line of demarcation.

Background

Backbone Architecture:

The choice of LAN backbone architecture was made based on a variety of factors including existing LAN infrastructure in place in the NCC, architecture of systems in use in the NCC (as in internal bus capabilities, etc.), traffic and bandwidth measurements made from the existing NCC Operational and Development environments, projections of traffic and bandwidth requirements for NCC98, and from research, experience, and partial evaluation of various types of LAN based architecture including shared-bus, collision avoidance/detection, duplexed, token passing, or switched architectures, some of which can even be combined. This choice had to be made in tandem with the choice of LAN technology (Ethernet, Fast Ethernet, 100BaseVG-AnyLAN, FDDI, ATM or some combination) and whether they were deployed individually or in combination, all have varying advantages and disadvantages which must be considered.

Using the *Hardware Specifications for the Network Backbone Architecture* document (available on the NCC98 home page) as a baseline for network backbone hardware evaluation, high-end LAN switches were deemed to be the focus of our research and evaluation testing. It was clear early on in the project that LAN switching held much promise as the “state of the art” in LAN backbone support. LAN switches provide a migration path from networks based on shared bandwidth hubs and bridges. They have become so popular and versatile that vendors have focused their R&D budgets on them such that they now surpass routers in terms of market share but more importantly now rival or surpass routers in terms of capability and performance. The architectural differences between routers and switches is becoming increasingly blurred with every passing day as more and more vendors push the limits by offering full routing capabilities with the performance of switching. As the premiere LAN backbone workhorse, the router has seen its day come and go. Routers are being relegated to the primary task which they were intended for, wide area network (WAN) support with a focus on external connectivity and security.

LAN Technology:

The choice of LAN technology also was made based on a variety of factors including existing LAN technology in place in the NCC, architecture of systems in use in the NCC (as in internal bus capabilities, etc.), traffic and bandwidth measurements made from the existing NCC Operational and Development environments, projections of traffic and bandwidth requirements for NCC98, and from research, experience, and partial evaluation of various types of LAN technology including Ethernet, Fast Ethernet, 100BaseVG-AnyLAN, FDDI, and even ATM.

Fast Ethernet is actually considered a generic term which has come to represent several types of high speed Ethernet with an effective data transfer rate of 100 Mbps including 100BaseT, 100BaseTX, 100BaseT4, 100BaseFL, 100BaseFX, and 100BaseVG. The terms 100BaseT and 100BaseTX are used interchangeably as are 100BaseFL and 100BaseFX (the TX and FX designations are the actual IEEE 802.3u physical layer device (PHY) wiring specs for duplex capability). 100BaseT4 is really just 100BaseT running over 4 pairs of wiring rather than two and uses a different signaling method. The differences between the 100BaseT and 100BaseVG specs are vast and they are not compatible. 100BaseVG was developed by Hewlett Packard as an alternative Fast Ethernet technology and is defined by IEEE 802.12 as Ethernet transmission over 100 Mbps twisted pair. It uses the same signaling method as 100BaseT4 but that's where the similarities stop as it is electrically incompatible with 100BaseT4. It was designed to work with 10BaseT, 10BaseFX, and 10Base5 but is not compatible with any of the 100Base variants. 100BaseVG uses a protocol known as VGanyLAN which is very similar to Token Ring and does not support collision detection. Even though it was accepted by the IEEE and ratified, it never gained acceptance by the general market and has suffered from very little vendor support aside from HP themselves and as well as various incompatibility problems. It was only within the last fiscal year that HP has finally begun to ship their own 100BaseT Fast Ethernet products such as NICs (they were filling customer requests with third party solutions).

Asynchronous Transfer Mode (ATM) switching technology was all the buzz just a couple of years ago (and still is in some camps) as the end-all technology for both the LAN and WAN. Highly touted by the media and the ATM Forum as THE network technology of the future, it was easy to believe all the hype and fall in rank. After the buzz wore off and network managers have had time to absorb their decisions to implement ATM in the LAN either to the desktop or just in the LAN backbone, there has been some regrets and some second guessing. ATM is most widely implemented in WAN backbones in conjunction with SONET due to it's

generous bandwidth support from OC-1 (51.84 Mbps) to OC-48 (nearly 2.5 Gbps). While ATM in the WAN arena is still making large strides and deployment here continues to increase, it is in the LAN that ATM deployment has been re-evaluated as the solution for every new LAN upgrade.

With most legacy networks running Ethernet (75% of all installed networks), there must be a substantial requirement to justify the implementation of ATM to the desktop. One such example would be a movie or video production house where a lot of work is done on high-end graphical workstations developing video imaging special effects. These users would be working with data, voice, video, and graphical imaging which would easily qualify as a premiere example of requiring ATM to the desktop. ATM was designed for multimedia traffic where the traffic requirements are made up of some sort of mixture of data, voice, video, and graphical imaging. This may even be in a LAN backbone where several networks which use one or two media types are collapsed with traffic of the other media types. Ethernet, Fast Ethernet, and FDDI are examples of variable length frame based transport technologies versus ATM which uses fixed length cells to transport information. It is this vital difference which lends itself to ATM's ability to be able to dedicate bandwidth and thus provide quality of service (QOS) and prioritization of traffic as part of its distinctive capabilities. This concept is a good idea and is actually deployed in the Cabletron MMAC Plus switch (it can be run with connection oriented switching software and has a cell based architecture to support ATM) which is evaluated in this report. Thus you can have the best of both worlds, a high-end LAN switch which can handle users who need ATM and users who only require a frame based technology like Ethernet all in the same switch.

Midway into the evaluation period, performance projections from network development specialists who had experienced high LAN utilization rates at other networks which are similar to the expected design of NCC98, caused a broadening of the LAN technology evaluation. Based on these performance concerns and on the critical, high reliability nature of the NCC, FDDI (Fiber Distributed Data Interface) began to be studied for use on the critical systems such as the SPSR and NSM servers, firewall, and NPG for its superior reliability and redundancy features as well as performance aspects. FDDI was compared side by side with Fast Ethernet in terms of performance, reliability, and redundancy both in a switched topology and a shared topology. It was determined through numerous research reports and white papers that switched Fast Ethernet rivaled or even exceeded shared FDDI and if run in a full duplex mode, Fast Ethernet was clearly superior in terms of performance. However, switched FDDI which can also be run in a full duplex mode takes top honors and even surpasses ATM running at OC-3 rates (155 Mbps). The

problem with the latter FDDI choice is that port density is very low compared to Fast Ethernet and much more costly. Another problem with FDDI is that even though it is run in a dual homed capacity to provide two paths to the switch from the end-station, that dual attached NIC is really a single point of failure. Advantages of FDDI are that either fiber or copper (twisted pair) could be used but we already have shielded twisted pair cabling in place which will meet FastEthernet requirements and unfortunately Hewlett Packard does not currently support CDDI (the copper equivalent of FDDI) interfaces on its equipment. If FDDI were chosen, fiber would have to be procured and installed as well as test equipment which could support FDDI.

Vendors:

Vendors chosen for the high-end switch evaluation were members of the so-called “Big Four” of networking products manufacturers. This statement adheres to item 6 of the *Hardware Specifications for the Network Backbone Architecture* document which requires that the vendors must have a minimum annual revenue of \$1 billion (US). Again, this specification was added to help limit the potential obsolescence, lack of future support, or unavailability of replacement/spare parts due to a smaller manufacturer being acquired (via merger or buyout) or simply going under. This industry is volatile in general and companies are constantly looking to keep an edge over competitors with new technology even if they have to buy it (which typically happens). This not to say that there were not very good switches from smaller vendors worthy of evaluation at the time. There were at the time evaluation research began and still exist very good switches that could be competitive from vendors like Xylan, Newbridge, and Madge. As a “case in point”, during the evaluation phase Madge encountered severe earnings shortfalls which led to large layoffs within the company simply due to stiff competition. Newbridge acquired their frame based (as in Ethernet, FDDI vs cell based like ATM) switching savvy from UB Networks which they bought out earlier in the year and were primarily an ATM switch vendor prior to the purchase.

Of the five vendors which met the revenue criteria in January 1997 when our evaluations began in earnest, only 3Com, Cisco, and Cabletron had switches which were chosen to be tested in the Building 12 N2 Lab. Bay Networks was eliminated due to un-remedied problems with the network management software which is used on all of their switching platforms (see *DEVLAN Implementation* in the **Evaluation Methodology** section of this report). IBM was eliminated due to research of third party evaluation test reports which showed noncompliance with a number of the

specifications in the *Hardware Specifications for the Network Backbone Architecture* document.

The chosen vendors were contacted and local NASA sales/engineering reps arranged for the delivery of the high-end LAN switches which had been chosen for evaluation in the Building 12 N2 Lab. The Catalyst 5000 switch was chosen from Cisco, the LANplex 6000 (now called CoreBuilder 6000) switch was chosen from 3Com, and the SmartSwitch MMAC Plus 6-slot switch was chosen from Cabletron. All three of these high-end switches met the basic specifications described in the *Hardware Specifications for the Network Backbone Architecture* document. This helped ensure that like items were being compared to like items.

Evaluation Methodology

1) DEVLAN Implementation

Experience gained from the implementation of LAN switches in the Development LAN (DEVLAN) in October 1996 has provided invaluable insight into one vendor's LAN switch performance and what may best be described as "lack thereof". The implementation consisted of stackable, workgroup size LAN switches and hubs from Bay Networks, Inc. Modular in design and rack-mountable, they fit the "spread out" DEVLAN environment fairly well with units installed in Room 213, Room 262, and in the SDF located off-site in GreenTec I. The initial installation and checkout uncovered that approximately 6 % of hardware was defective out of the box. This fits the unofficial industry standard of between 5 % to 10% rate of defective hardware out of the box. Problems specific to the Bay equipment ranged from bad ports in a module to defective LEDs.

Cut-over of the DEVLAN nodes uncovered what was to become the "Achilles heel" of the Bay switches and hubs. Bay's network management software which is used for switch/hub management and configuration and must be used to perform any tasks associated with virtual LANs was highly problematic and bug infested. Even with a couple of visits from their Systems Engineers, we were unable to get satisfactory results and performance from either the Windows version or the Unix version. This network management problem was deemed unsatisfactory for eventual implementation in NCC98 and it was decided that other vendors products should be evaluated. A recent upgrade of the DEVLAN thoroughly highlighted Bay's poor performance with their network management software (Optivity) especially when contrasted to the network/switch management software from the recently evaluated vendors (Cisco, Cabletron, and 3Com).

Additional network infrastructure has been implemented in the DEVLAN which gives us feedback for consideration of potential usage in NCC98. Items such as transceivers, media converters, manageable and unmanageable shared LAN hubs, sniffers, and routers have been or are currently in use.

2) Building 12 Room N2 Test Lab

Testing of the high-end eval switches was performed in the Code 530.5 N2 Test Lab located in Building 12, Room N2. Additional networking device testing is also undertaken here. Several workstations from SUN Microsystems and Hewlett Packard are utilized as well as several Windows based PCs. Access to a Distributed Sniffer System from Network General is also provided. The vendor's switch management/configuration software was loaded and evaluated on these test nodes as well as several third party network management and analysis programs. Dumb terminals are likewise used for switch configuration purposes.

3) Third Party Test Reports

Third party test reports were utilized to make up for testing which could not be accomplished due to the limited resources in the N2 Test Lab. These reports were used to evaluate items like switch packet throughput, multiple port performance capabilities (e.g., latency measured through the switch when 12 nodes are making simultaneous connections through the switch which are equal volume connections), and other various performance related measurements. This type of testing requires extensive test and measurement equipment such as traffic generators with multiple interfaces, network analyzers with multiple interfaces to various LAN technologies, and enough nodes to support the required configurations with the appropriate interfaces. Research labs and consulting firms which specialize in this type of testing are the pre-eminent Harvard Network Device Test Lab (headed by Scott Bradner), Strategic Networks Consulting, Inc., The Tolly Group, Inc., and Decisys, Inc. to mention just a few which were used in the high-end switch evaluation process. Additional evaluations and test reports from network related trade periodicals were used also which often times provide guidelines on what to look for when performing your own evaluations.

4) Interviews with Current Users

Interviews were made with network consultants, network designers, system administrators, network engineers, and users of various networks both on-site at GSFC and off-site employing network equipment (switches, hubs, routers) from the three vendors whose switches were under evaluation. Some even used the same switches which we had under evaluation. This provided additional insight into problems experienced in real on-line configurations and the respective vendor's support.

5) Establishment of Criteria

The criteria which were established to base there commendation for the high-end switch on are as follows:

- 1) Reliability (based on MTBF figures and feedback from existing implementations)
- 2) Redundancy (must minimize or eliminate single points of failure)
- 3) Architecture/Performance (includes items such as switching fabric, aggregate backplane capacity, throughput, latency, packet loss, buffering, congestion design/measurements, et al)
- 4) Cost (on a per port basis)
- 5) Administration (includes configuration and management)
- 6) Maturity/Life Expectancy of the product/family/product line
- 7) Features/Scalability (required and distinguishing/port density and LAN technology)
- 8) Vendor stability (specifically within the switching arena)
- 9) Local support (both pre and post sales)
- 10) User implementation experiences and satisfaction

It would be very hard to rank these criteria in order of importance and the order they are listed does not necessarily infer this, however, it was decided by previous discussions with the customer that reliability, redundancy, and cost were probably at the top of the list, with architecture/performance, administration/management, and features/scalability forming a second priority grouping.. The remaining criteria would then form the third priority grouping.

Evaluation Data

N2 Lab Testing:

1) Lab Testing Goals

The testing in the lab was primarily to determine ease of use in the set-up and configuration of the high-end switches, the same for the switches' network management software, and the thorough checkout of switch features intended for use in the NCC98 design.

2) Set-up and Configuration

The evaluation switches were arranged for via local or NASA vendor reps. Some were brought in by the local vendor reps themselves and some were shipped. Cisco and Cabletron sent their switches by shipping carriers. 3Com delivered their demo switch in person accompanied by a network engineer who set up and attempted to configure the switch during his visit. Each vendor was asked to provide a minimum of 4 Fast Ethernet ports and 8 Ethernet ports, all necessary chassis, power supplies, fans, management/monitor modules, and all necessary network management software to be able to fully configure all of the switch's capabilities. FDDI concentrator and switch modules were also received at a later date during the brief FDDI evaluation period mentioned earlier in this report. Up to date switch images (software which runs internally in the switch modules themselves) were also requested either pre-loaded or via some form of transferable media.

Cabletron: Cabletron's switch (SmartSwitch MMAC Plus 6-slot) came boxed and packed in its original shipping containers and it was evident that the equipment was brand new out of the box. Set-up was straightforward and relatively simple. The modular power supplies were installed first; they inserted into the front of the chassis for easy accessibility. Next the fan module and system monitor modules were installed in their respective positions (also in the front for easy access). A 12 port Auto-Negotiating 10/100 Fast Ethernet module was inserted into slot 1 and a 36 port Ethernet module was inserted into slot 2. Initial configuration was a simple task of using a terminal connection attached to one of the COM ports on the System Monitor module and turning on one of the redundant power supplies. The switch utilizes a distributed 48-volt DC power system similar to those used in high reliability super computers and telco equipment. It is a power hungry unit and each supply provides

up to 1000 Watts of power in a load sharing, fully redundant configuration. Upon boot-up the administrator is presented a password protected, menu driven configuration and management platform by which to set-up IP addresses, port configuration details such as data rate, full or half duplex, etc.

Cisco: Cisco provided a demo switch (Catalyst 5000 aka CAT 5000) which was poorly packaged for the trip via FEDEX (very little insulation material which did not fit the box correctly). As a result, it was damaged upon arrival, suffering a bent power supply insertion handle and generally suffering some cosmetic damage. However, it seemed to function normally upon initial boot-up and subsequent testing. The switch was shipped with the requested configuration of Ethernet (24 port Ethernet module) and Fast Ethernet (12 port 10/100 Auto-Negotiating Fast Ethernet module) in addition to the required Supervisor Engine module. The Ethernet module seemed to occasionally exhibit abnormal LED indications which were never fully resolved. Set-up was already completed as the switch came with all its modules pre-inserted since it was a demo which basically goes from one customer eval to the next. The modules were removed and re-inserted to gain a feel for this exercise. All were front loaded for easy access. The power supplies are also 48 volt DC, fully redundant, load sharing and provide up to 376 Watts individually. Configuration via a terminal based serial connection is possible but requires extensive knowledge of Cisco's command-line interface which has been borrowed from their router centric configuration philosophy. This type of interface is really unnecessary for configuration of a high-end switch as it takes the level of competence required to configure and maintain the switch to a higher level. Although with the recent advent of Layer 3 switching and routing added to switches, the preceding statement is losing its validity. Once you can get an IP address into the switch, the remaining configuration can be completed via the SNMP based switch management software (CiscoWorks) which is a much easier process instead of having to type "<command> ?" repeatedly to get through the command-line interface structure (maybe this is what the 5000 stands for in CAT 5000). VLANs must be created and administered from a different switch management software package called VLAN Director.

3Com: The 3Com switch was a 4 slot chassis (LANplex 6004) but architecturally (relative to the switching architecture as opposed to physical packaging) identical to the 12-slot chassis (LANplex 6012) which would be required for the NCC98 installation. The 4-slot chassis only has room for a single power supply which isn't modular and therefore it cannot be considered for implementation (aside from the fact that it could not meet the port density requirements - it runs out of slots). The requested configuration was met over time (the 8-port Fast Ethernet module took

months to obtain due to its recent release and filling of advance orders) with a management module and a 16-port Ethernet module which included an uplink port for Fast Ethernet. The 12-slot chassis includes modular 1000 Watt redundant power supplies. Although some of the 3Com switch modules were demo equipment, some were brand new out of the box. The 3Com rep demonstrated his apparent lack of knowledge on the switch set-up by installing the management module in the wrong slot. This didn't matter since the management module requires an odd choice of serial COM for initial configuration via a dumb terminal. A 2 pair cable using RJ-11 connections is required which neither we in the lab nor the 3Com rep seemed to have handy. Although he redeemed himself fairly quickly by shipping one out the next day via FEDEX (however, this too was the wrong cable and we eventually received the correct one the following day also by FEDEX). Initial configuration was a little more interesting with the LANplex due to odd numbering schemes for the module slots and for the module ports also. The port numbers on the chassis don't match the port numbers depicted in the terminal set-up program which was very confusing at first. Once this was understood, configuration of basic items was fairly straightforward. They even provide the option of using a standard menu based application or a command line interface for config/admin chores.

3) Switch Management Software

The problems with Bay Networks network/switch management software (both Windows and Unix versions) made this an item to key on with the other vendors. It can be stated that none of the three network/switch management programs (Unix versions) looked at from the three eval vendors (3Com, Cisco, Cabletron) came anywhere close to the problems experienced with the Bay Networks software. Ranked by order in ease of use and least amount of minor bugs, Cabletron led followed by Cisco. We were not able to get a working Unix version of 3Com's Transcend network management software installed as resource problems stalled our plans. This package requires HP OpenView for UX which we did not get installed by the time that we had decided not to expend further test time on the 3Com evaluation due to deficiencies in meeting the established hardware criteria. However, a three hour seminar was attended which provided an overview of 3Com's Transcend network management software and the impression gained from this experience was that Transcend was a solid product which integrated tightly with HP OpenView but that its VLAN strategy was geared for usage with ATM although totally functional for use with frame based LANs. Windows versions of Cabletron's VLAN management software (for NT 3.5) and Cisco's CiscoWorks management software (for 95) were also evaluated early on and had significantly more problems. The bottom line here is that switch/network management software is best run under a high-end platform such as Unix based

workstations. One caveat to Cabletron's stand alone VLAN Manager program (Unix version) is that after the next release (1.6) it will no longer be supported on HP UX. As a result, this program would require a SUN Solaris workstation (recommended) as a platform for the switch management software.

4) Switch Features and Functionality

Once the switches were fully set-up and configured, testing was performed to evaluate the features which make these switches so called "high-end" switches. These are items such as support for virtual LANs (VLANs), network analysis provisions, port density based on LAN technology (Ethernet, Fast Ethernet, FDDI, ATM, etc.), Layer 3 switching capability, full duplex capability, fault tolerant capabilities, packet filtering capabilities, and any distinguishing features that set the switch apart from the competition. The following tables summarize and compare features found in each of the three switches which were evaluated (Note: these are the features which were available at the time the evaluations took place...vendors are constantly adding new features and upgrades which may not have existed at the time of the evaluation of each particular model). The tables are broken down into features that were required by the *Hardware Specifications for the Network Backbone Architecture* document versus features which were requested. It should be noted that some of the requested features have now become required based on the *Operations Concepts* document (they appear shaded with cross-hatch).

Cabletron: The MMAC Plus features definitely are a step ahead of the competition. Switch architecture is a big difference. Distributed management with no management module as a single point of failure is a big difference compared to the other two. Both of the other two switches have management modules as single points of failure which is even worse for Cisco since they do their actual switching on the management module. Cabletron also has redundant active backplane structures which only 3Com comes partially close to duplicating (3 FDDI rings). Additionally, a unique method of using time division multiplexing in the ASIC based switch fabric makes congestion issues meaningless as well as provides fully non-blocking bandwidth capacity. Redundancy is provided in every possible element of this switch with multiple backplane structures, dual power supplies, etc. which make it the closest of the three switches to achieve complete fault tolerance. A big highlight is the method for supporting redundant links between switches (trunks). It is known as active mesh topology which means that all trunks between switches are active trunks carrying traffic and are thus load sharing by default. This is achieved by an OSPF (Open Shortest Path First) based routing protocol known as Virtual Link State Protocol (VLSP) which has been adapted for Layer 2 usage. Contrasted with the other two switches in this area, active

mesh VLSP is clearly superior to the Spanning Tree Protocol (STP) which was adopted by the IEEE and is specified in IEEE 802.1d in the early days when bridges were developed. STP sets up redundant standby links which are only activated when the link or trunk designated as primary fails. STP takes 30 to 50 seconds to re-converge during which time, no traffic is being forwarded through the networking devices running it. VLSP re-convergence is measured in milliseconds.

Table 1: Switch Features Required

Vendor	Cabletron	Cisco	3Com
Model	MMAC Plus 6-slot	Catalyst 5000	LANplex 6012
Features Required			
Modular Chassis	Yes	Yes	Yes
Duplex Capability	Yes	Yes	Yes
MAC Addr per port (>=1024)	16000	30000	8132
Hot swappable modules	Yes	Yes	Yes
Any single points of failure	None known or found	Bus structure, Supervisor engine	Management module and bus, HSI bus
Support Redundant Links	Yes Uses Active Mesh	Yes Uses 802.1d	Yes Uses 802.1d
Virtual LANs	Yes	Yes	Yes
SNMP Manageability by Ind Std Apps	Yes	Yes	Yes
SNMP Management Required for VLAN Admin (Vendor)	VLANManager	VLANDirector	Transcend
In band/Out of band Mgmt Ports	2 Serial COM, 1 Side Band Ethernet	RS232, 2 Fast Ethernet, 2 MII, 100BaseFX	2 Serial COM, 1 Ethernet, RS232, FDDI
Support for external network analysis	Port Mirroring, Call Tapping	ESPAN	Roving Analysis
MIBs Supported	MIB II, Ethernet, Bridge, FDDI SMT, LEC, ATOM, Cabletron	MIB II, Ethernet, Bridge, FDDI SMT, LEC, ATOM, ILMI, Cisco	MIB II, Ethernet, Bridge, FDDI SMT, LEC, ATOM, LANplex
RMON Support	Yes	Yes	Yes

Table 2: Switch Features Requested

Vendor	Cabletron	Cisco	3Com
Model	MMAC Plus 6-slot	Catalyst 5000	LANplex 6012
Features Requested			
Switch Architecture			
Management	Distributed	Centralized	Centralized
Internal Data Type	Frame and Cell based	Frame based	Frame based
Backplane Structures	Dual FDDI rings, Dual 64 bit internal, Cell Transfer Matrix, Dual System Mgmt	Single <i>Data Switching Bus</i>	3 FDDI rings, <i>High Speed Interconnect</i> , System Mgmt
Fabric Bandwidth	Non-blocking	Partially Non-blocking	Non-blocking
Aggregate Bandwidth	5 Gbps	1.2 Gbps	19.5 Gbps
Implementation Type	ASIC for switching, RISC for management	ASIC and RISC combination	ASIC for switching, RISC for management
Congestion Control	Time Division Multiplexing	Netflow Switching	Elastic Packet Buffering
Buffering per Port	2 MB Dynamic + 64KB	192 KB	Dynamic 64KB to 2MB
Technology Life Cycle	5-10 Years	3- 5 Years	2- 3 Years
Broadcast Control	Yes	No	No
Packet filtering	No	No	Yes
Multi VLAN per port	Yes	Partial	No
Routing Protocols	No	RIP, OSPF	RIP
Layer 3 Switching	Yes	No	Yes
Multicasting Support	No	Cisco Gateway Management Protocol	Distance Vector Multicast Routing Protocol

Cisco: The CAT 5000 and the entire Catalyst Switch family was bought out from a company called Kalpana a couple of years ago. Cisco, realizing they were late to market with any type of competitive switching product, decided to go after this expected huge market for LAN switches by choosing a good design from a startup. From a revenue standpoint, the choice was more than they could have hoped for as they now lead the LAN switch market in this area, with 3Com, Cabletron, and Bay right on their heels. Even more impressive is the fact that they made more revenue from switches than routers in fiscal year 1996. However, from a architecture and performance standpoint, they realize they are somewhat behind the curve and are probably not considered in the top echelon based on most third party test reports and trade magazines. Although what they lack in competitive switching performance and capability, they make up for in marketing muscle. Limitations within the switch's architecture inhibit them from entering the upper echelon and they may never unless they market a next generation switch based on a refined architecture (which they are probably secretly working on already). Their latest entry, the Cat 5500 is really just a larger Cat 5000 providing more slots in order to overcome some of the deficiencies occurring in the Cat 5000 chassis such as the supervisor engine which apparently had a high failure rate and was a single point of failure. The Cat 5500 uses two supervisor engines (management modules) to overcome this limitation as well as a lot of re-design to incorporate new features (layer 3 switching, multiple active trunk links, automatic VLAN configuration) already stable and mature in the upper echelon switches.

3Com: The LANplex switch provides an architecture designed with elements of brilliance of which the only shortcomings were it's lack of foresight towards scalability into additional technologies. One area where this is apparent is the lack of a cell based data format for compatibility with ATM. The other more disappointing problem is the odd way in which the Ethernet modules tie into the switch's backplane structure which may be the result of the switch's evolution into higher speed and bandwidth support. The problem is that all of the 10BaseT Ethernet modules tie into the FDDI backplane whereas the Fast Ethernet module and FDDI Switching module tie into the High Speed Interconnect bus (the 19.2 Gbps bus) and therefore any connections between a node on the Ethernet card and another node on the Fast Ethernet card must be translated (this adds latency and unnecessary overhead) into FDDI frames and then back to Ethernet again (thus the conspicuous absence of this particular switch from most third party test reports...especially any dealing with performance measurements and latency). What's even worse is that in order for the switch to make this translation it must be done via the FDDI Switching Module, an unnecessary additional module costing \$8000 and occupying an additional chassis slot to do nothing more than switch

Ethernet frames to Fast Ethernet. This is considered a critical architectural faux pas.

MTBF Data:

Rather than go into an explicit chart depicting each of the three switches and their respective components' associated Mean Time Between Failure (MTBF) figures, the following worst case scenario availability calculation is provided using the lowest MTBF figure (51,000 hours calculated using Bellcore specifications for an Ethernet switch module) available from any of the figures provided by the vendors on every switch component which would be required to satisfy the intended implementation in NCC98.

$$\text{Availability} = \frac{\text{MTBF}}{\text{MTBF} + \text{MDT}}$$

where MDT = Mean Down Time ~ which is defined as the summation of the time required to acknowledge the event, the time required for maintenance personnel to get to the component at fault, the time required to diagnose the problem, and the time required to replace the faulty component or module.

For this example, assume it takes one hour to acknowledge the event(via NSM or the switch management software), get maintenance to the appropriate switch, determine which module is at fault, replace the module, and reconnect the patch cables. This is the mean down time or MDT figure.

$$\text{Availability} = \frac{51000}{51000 + 1.0} = 0.99998039 = 99.99804 \%$$

which is well within the limits for 99.990% availability for the backbone hardware as specified by the *Hardware Specifications for the NCC98 Network Backbone Architecture* document. This scenario doesn't affect system availability since overall system down time would not occur due to redundant switches in OPS. Actually, based on the MTBF figure used above, network administration and maintenance personnel would actually have nearly 5 and 1/2 hours of MDT to work in if they were to meet the 99.990 % backbone availability requirement and up to 10 hours to meet the 99.980 % system availability requirement (assuming actual system downtime occurred).

Third Party Test Reports:

Third party test reports were used to substitute for performance testing which could not be accomplished due to limited resources. Several consulting groups and trade magazines perform these services and publish the results free of charge (usually via the web) or in monthly articles. Absence of a particular vendor's product is usually seen as an admission of weakness in the particular area that the test will concentrate on. Of course, all test results must be thoroughly scrutinized based on test methodology and goals. Consistency among various tests is usually a good way to compare results.

Cabletron: Cabletron's MMAC switches running SecureFast images (an image is firmware code that resides in FLASH or NVRAM) fared very well in third party test reports within the last year. Notable reports for this particular switch were:

- 1) Networld+Interop IntraNet Excellence Award in Infrastructure category (1997)
 - 2) Rated "Superior" in the categories of VLAN, RMON support by The Tolly Group (independent consultants) (1997)
 - 3) Consistently rank highly when tested by Strategic Networks Consulting, Inc. in conjunction with the Harvard Network Device Test Lab under the direction of Scott Bradner in performance, features, and functionality. (1996/97)
 - 4) Scored highest (scored an A or A- in every category) in Communications Week magazine test of VLAN strategies(2/97)
 - 5) Took "Tester's Choice" award for Virtual LANs from DataComm magazine (5/97)
- (grade: A)

Cisco: Cisco is very selective about which third party tests they enter. They are very shy about VLAN competitions and no reports or tests could be found which covered their VLAN strategy. The only award within the past year for the Catalyst 5000 was Datamation magazine's "Product of the Year" award.(2/97) (Source: Cisco Systems Product Family Overview documentation)Although this award was deemed to be of little more than a "product capabilities" type of review and did not actually put the switch through any actual measurable testing. (grade: C)

3Com: 3Com's LANplex 6000 was also conspicuously absent from any third party test reports primarily due to the aging architecture of the 6000. Several reports were found on the LANplex 2500 which is a smaller version of the 6000 and more aptly considered a mid-level switch rather than a high-end switch. (grade: C)

User Interview Highlights:

Cabletron: Cabletron switch users were fairly abundant, but harder to find were users of the actual MMAC switch family. The only users that were accessible via an AT&T contact met at a training class who said that they had been installing a large number of the MMAC switches at the time and that they were working well with no identifiable concerns or problems. More recently, the CNE has purchased three of the MMAC Plus 6-slots for use in the CNE backbone. The only discernible complication during installation was a single bad Ethernet port on one module. This was far below the 5% to 10% typical defective "out of the box" ratio referred to earlier in this report. (grade: A)

Cisco: Users of the CAT 5000 who were accessible for interviewing were plentiful. Experiences of users from within Nascom and the CNE here at GSFC have been mixed. The entire HST network backbone is built upon Cisco devices with Lightstream 1010 and the new CAT 5500's forming the core of the backbone. CAT 5000's are used merely as edge devices. The enormity of the HST implementation provided an excellent source for feedback. In general, the network has been working great due to the high degree of reliability designed in (lots of switches for alternate paths in the backbone) and due to the abundance of raw bandwidth to all points (OC-3 backbone). The CAT 5000 taken individually had some early problems, some of which have been remedied to date. Specifically, the power supplies had a high failure rate early on until redesigned, hot-swappability of modules is occasionally flaky, and the supervisor engines are a single point of failure, thus when they fail, the switch is inoperable. Cisco realized how critical the later problem was and that's why they redesigned the supervisor engine (ironically called the Supervisor Engine II) as well as provide for two supervisor engines in the CAT 5500. While they were at it, they provided full non-blocking switching support for the portions which did not yet support it too. This new redesigned module was not available at the time of evaluation. Nascom feedback was based on a single CAT 5000 borrowed from HST which proved problematic to configure VLANs with. CNE has many CAT 5000's and their usage is similar to HST's, they are used primarily as edge devices and not as backbone switches. They too have experienced hot swap inconsistency and are aware of the supervisor engine weakness. (grade: C)

3Com: Other than routers and low-end networking products such as NICs and hubs, 3Com switch users were found to be in short supply for interview purposes. The high cost of these switches compared to the competition probably has more to do with this than any other factor. (grade: N/A)

Vendor Support:

Cabletron: Cabletron's support was handled out of Rochester, NH by the Federal Accounts Manager assigned to NASA accounts. Local support was provided by the Federal Sales Office in Herndon. The Federal Accounts Manager assigned to NASA accounts has changed three times during this evaluation period, however, the level of support has remained consistently very good. Phone calls and email have been handled expediently. Quotes, documentation, and reference material were adequately and expediently prepared and provided. Requests for technical help have always been met with technical competence, however, not as expediently as we would have liked. Overall Cabletron's support was considered very good. (grade: B+)

Cisco: Cisco's support representative changed early in the project, so it was not possible to make any judgement at this point. The rep who took over works out of the Herndon office and is the Federal Systems Account Manager assigned to NASA accounts. The level of support received at this point was excellent with calls and email handled expediently. Quotes, documentation, and reference material were adequately and expediently prepared and provided. Requests for technical help were always replied to with superior technical competence. Technical issues which were discussed were addressed with complete candor in terms of what could and couldn't be done with the vendor's product. Overall Cisco's support was considered excellent. (grade: A)

3Com: 3Com's support was handled out of the Vienna office by the Federal Territory Manager assigned to NASA accounts. Aside from the set-up and installation miscues on the first day (which implanted a bad impression of expected technical support), support was generally good with calls and email handled expediently. Quotes, documentation, and reference material were adequately and expediently prepared and provided. They were the slowest to respond to the initial request for evaluation gear, although they responded better when additional switch modules were asked for. However, even this order had problems as additional units were sent out which were not requested. These then had to be turned around and sent back which wasted time, so overall 3Com's support was considered satisfactory. (grade: C)

Cost Data:

Cost data is based on the required configuration to support all of NCC98 including the T&T, ANCC, and OPS with a switched Ethernet and switched FastEthernet topology employing redundant trunk links between four switches. At the time of this report, the configuration was estimated to be 8 ports of Fast Ethernet and 21 ports of Ethernet per switch including on board expansion of 30%. This excludes additional infrastructure such as cabling (of which most of the CAT 5 shielded twisted pair is already in place), shared LAN hubs, NICs on any nodes, and any other ancillary equipment. For the component level cost matrix (based on SEWP II prices in conjunction with vendor quotes as appropriate) see the appendix. The bottom line for each of the three switches evaluated appear below.

Vendor	Cabletron	Cisco	3Com
Model	MMAC + 6-slot	Catalyst 5000	LANplex 6012
Cost Per Port (Actual)*	\$ 676	\$ 975	\$ 1525
Port Config:			
Ethernet	36	24	32
Fast Ethernet	12	12	8
Total Cost	\$ 129,770	\$ 140,520	\$ 244,085

*Cost per port (Actual) data are based on the actual number of *operational traffic capable* data ports in the switch after the required number of ports have been met. To clarify, even though a switch module could be purchased which meets the minimum requirement, a module with a higher port density could be substituted for the same price thereby decreasing this cost figure.

Recommendations

NCC98 Backbone Design :

The following points are recommendations for the NCC98 backbone design in general and are complementary to the choice of high-end LAN switch which follows.

- 1) **Cabling:** Due to the switched Ethernet and switched Fast Ethernet (100BaseTX) topology, Category 5 twisted pair cabling is required for all links as well as components meeting Category 5 standards such as connectors, jumpers, and patchpanels. This is especially important for the Fast Ethernet links as the 125 MHz signaling rate presses the limits of copper based twisted pair cabling and components. Installation of any new Cat 5 cabling runs should also follow strict adherence to EIA/TIA 568/569 standards which specify details such as minimum bend radius (shall not exceed 4 times outer diameter of jacketed cable), maximum pull force, maximum cable length, avoidance of kinks, etc. Problems resulting from improperly installed cables can result in impedance mismatches, crosstalk, near-end crosstalk (NEXT), and increased attenuation all of which can result in transmission problems for Fast Ethernet signals. All pre-existing cable runs, whether they are to be used for Ethernet or Fast Ethernet, should be inspected for any problems or conditions not meeting the standards referred to above. This way a cable run could support Fast Ethernet in case requirements justify its use.

It may be necessary to use 62.5/125 micrometer multi-mode fiber optic cabling for runs from Room C-130 and Room 262 in Building 13 to the ANCC (Building 3, Room S60) which are the Fast Ethernet trunks between the switches if the distance is found to exceed 100 meters (the specified distance limitation). Also, 50/100 micrometer multi-mode fiber can be substituted if necessary. These cable runs would follow 100BaseFX Fast Ethernet specs.

- 2) **Dual Fast Ethernet (100BaseTX) NICs on K-Servers:** Servers with the potential for high traffic volume should be fitted with dual Fast Ethernet NICs. Each of the two interfaces will be dual homed to each of the two switches placed in OPS. This is recommended for the K servers so that they will not have to use a switch trunk were one of their links to fail if they were only single homed alternately (i.e., for efficiency).
- 3) **Fast Ethernet(100BaseTX) Backbone:** Fast Ethernet (100BaseTX/FX) is recommended for use in the backbone as trunks between switches. By running the Cabletron switch's SecureFast images, an active mesh

topology will automatically be enabled and provide the high reliability, high efficiency backbone discussed earlier in the report. Further increasing efficiency and backbone bandwidth capacity, the trunks can be run in a full duplex mode thereby increasing the effective backbone bandwidth to 200 Mbps between switches.

- 4) **Side-Band Switch Management:** Cabletron's VLAN Manager switch software should be run on a SUN Solaris platform (2.5 or better) and tied to a shared LAN hub with connections to each of the four switches via their side-band Ethernet (i.e., not on the same backplane as the data traffic) connection which is present on each switches System Monitor module. This has the twofold effect of providing internal/external security of the switch management software as well as preventing any general network related problems from interfering with switch/VLAN management duties. A second on-line and fully active (running the VLANManager switch software) SUN Solaris node (2.5 or better) can be provided and also tied to the shared LAN hub in order to provide redundancy and still meet the backbone availability requirements. Multiple active VLANManager clients can be on-line at anytime. If the prime node were to go down, the VLANManager server would be activated on the redundant machine and the system would be fully manageable after a few seconds.
- 5) **Management Based VLAN:** It is recommended that network system management (NSM) functions be separated from operational traffic with the use of VLAN technology as much as is feasible. Specifically, SNMP manageable hubs are recommended for use for the external heartbeat LAN used for the K servers, the switch management LAN (described above), and the external firewall IONET connections. The management ports of the hubs (Ethernet based) should be connected to the appropriate switch ports which would be designated part of a *secure* "Management" VLAN. A *secure* VLAN is one in which no connectivity can be established with members inside the VLAN from the outside without the use of a layer 3 routing protocol. The internal (via the switches) heartbeat LAN used for the K servers should also be placed on a *secure* VLAN in order to establish some form of isolation.
- 6) **Dynamic Routing Protocol for IONET Connections:** A dynamic routing protocol, specifically OSPF, should be used for the external IONET connections in order to utilize the dynamic backup features of the protocol. This would enable the routers to automatically re-converge traffic to a backup operational data link should a failure occur on the primary active link. This recommendation is in lieu of the use of static routes currently being used for NCC97. This is currently being worked on by the NCC97 Development Team in coordination with Nascom IP Transition engineers as it requires changes on both sides of the links.

High-End LAN Switch:

Cost, reliability, and redundancy were considered to be the first priority for choosing the high-end LAN switch. The second priority group was scalability, management, and features/functionality/performance. Lastly but still of import were user satisfaction/experiences, vendor support, and life expectancy of the switch architecture. As far as the first priority group, the Cabletron switch is the leader in all three areas of reliability, redundancy, and cost.. The 3Com switch would be considered a second choice here only due to the fact that reliability and redundancy probably should outweigh cost and 3Com was judged to have the edge over Cisco in this respect due to the CAT 5000's use of the supervisor engine as a central switching module versus the LANplex's distributed switching per module approach. Even though Cisco expects to correct this and may have recently, it is still a new design which has yet to be field-proven and thus become mature and stable.

As far as the second priority grouping, again Cabletron is judged to be the best with Cisco taking second honors above 3Com. The choice for second again was a tough one but Cisco wins out primarily due to the fact that it provides increased scalability over the LANplex (3Com) and is attempting to draw closer to the upper echelon of high-end switch manufacturers with it's rash of re-designs within the last year which have added or are expected to add a variety of new features and functionality (e.g., true layer 3 switching, active trunk links, increased reliability and redundancy). The only drawback to re-design and new features are the time it takes for them to become mature and stable (i.e., usable).

Lastly, the third priority grouping again places Cabletron first primarily based on the life expectancy of the switch (5 to 10 years) and it's architecture when compared with the other two switches. 3Com's local NASA rep made it clear that the various architectural problems of the LANplex 6000 have limited it's life expectancy to 2 to 3 years at best. Cisco continues to squeeze new life out of its market leading Catalyst switch family through re-designs and re-works as evident with the introduction recently of the CAT 5500, a 13-slot chassis which offers increased scalability, reliability, redundancy, and attempts to overcome the performance shortfalls of its predecessor, the CAT 5000.

Cabletron's switch comes out on top in almost every category of the evaluation and when looked at from the overall standpoint, is the best switch for NCC98 and beyond. The **Cabletron** switch (MMAC Plus 6 slot) is recommended as the clear choice among the three switches evaluated for NCC98. The choice from Cisco would be recommended as the second overall selection and the choice from 3Com would be recommended only as a last resort.

Sources

- Axner, David, "Cabletron versus Cisco: Evaluating Network Architectures", Telecommunications, February, 1997.
- Boteler, Michael J., BISnet Platform Manager, AT&T Information Technology Services, Herndon, VA.
- Cisco Systems, Inc., *Scaling the Campus Intranet*, San Jose, CA, 1997.
- Freeman, John and David Passmore, *The Virtual LAN Technology Report*, Decisys, Inc., Sterling, VA, 1996.
- Kumar, Siva S. and David Newman, "Ethernet Switches: Quantity, Not Commodity", Data Communications, November, 1996.
- Mascari, Michele, Communications Specialist, Nascom ATM/NILE Lab, Lockheed Martin Space Mission Systems & Services, Seabrook, MD.
- Metz, Christopher Y. and George C. Sackett, *ATM and Multiprotocol Networking*, McGraw-Hill, 1997.
- Nemzow, Martin, *FastEthernet Implementation and Migration Solutions*, McGraw-Hill, 1997.
- Rash, Wayne Jr., "The State of the VLAN", Communications Week, February 3, 1997.
- Wentz, Tom, Federal Account Executive, Madge Networks, Inc., McLean, VA.
- Wharton, Dan, Network Design Consultant, HST Backbone Project, Greenbelt, MD.

WWW Sites:

www.baynetworks.com	Bay Networks, Inc.
www.cabletron.com	Cabletron, Inc.
www.cisco.com	Cisco Systems, Inc.
www.3com.com	3Com, Inc.
www.snci.com	Strategic Networks Consulting, Inc.
www.tolly.com	Tolly Group, Inc.
www.decisys.com	Decisys, Inc.
www.newbridge.com	Newbridge Networks Corp, Inc.
www.madge.com	Madge Networks, Inc.

Appendix

Eval Switch Cost Matrix

	LANplex 6012	MMAC Plus 6-slot	CAT 5000
Chassis, Fans, etc.	19584	2472	3141
Two PS	included	2800	925
Mgmt Module	included	777	5176
10BaseT Mod (#ports)	25920 [32]*	13197 [36]**	3696 [24]
100BaseTMod (#ports)	7096 [8]	13197 [12]	7396 [12]
Additional Modules Required	8520	N/A	14796
TOTAL COST	61120	32443	35130
Cost/Port	1528	676	975
* = Requires two modules, 16 ports each to meet requirements			
** = Only 24 port module is required but can get 36 port module for the same price			